

Informed multi-objective decision-making in environmental management using Pareto optimality

Maureen C. Kennedy^{1*}, E. David Ford², Peter Singleton³, Mark Finney⁴ and James K. Agee²

¹University of Washington, Quantitative Ecology and Resource Management, Box 352100, Seattle, WA 98195–2100, USA;

²University of Washington, College of Forest Resources, Box 352100, Seattle, WA 98195–2100, USA; ³USFS Pacific North-west Research Station, 1133 N. Western Ave, Wenatchee WA 98801, USA; and ⁴Rocky Mountain Research Station, Fire Sciences Laboratory, 5775 Hwy 10 West, Missoula MT 59801, USA

Summary

1. Effective decision-making in environmental management requires the consideration of multiple objectives that may conflict. Common optimization methods use weights on the multiple objectives to aggregate them into a single value, neglecting valuable insight into the relationships among the objectives in the management problem.

2. We present a multi-objective optimization procedure that approximates the non-dominated Pareto frontier without the use of weightings, allowing for visualization of the trade-offs among objectives. The non-dominated Pareto frontier is approximated by the simultaneous optimization of a vector objective function; two vector objective functions are defined as non-dominated if improvement with respect to one objective is at the detriment of another objective.

3. We demonstrate the method with a case study for the optimum distribution of forest fuels treatments that reduce the impact of fire on a forest. The multiple objectives are to protect habitat of an endangered species, protect late successional forest reserves and minimize the total area treated. In the comparison of three optimization searches, the number of non-dominated solutions increases with the dimensions of the objective space, but with only two objectives the search is ineffective in minimizing fire impact in the different landscape types. Key challenges include the extensive computation time required to approximate the non-dominated set, and reducing the number of solutions that are analysed in detail.

4. *Synthesis and applications.* The multi-objective optimization program presented can be adapted to other environmental management problems, and easily incorporates a wide range of quantifiable objectives. This tool provides decision-makers with a set of alternatives that estimates the full range of trade-offs among multiple objectives and provides a common ground from which dialogue can come to an informed compromise and decision in environmental management problems.

Key-words: fire management, fuels management, multiple objectives, non-dominance, optimization, Pareto frontier, spatial structure

Introduction

MULTI-OBJECTIVE OPTIMIZATION AND ENVIRONMENTAL MANAGEMENT

Multiple criteria analysis and multi-objective optimization have been utilized to design decision support systems for a variety of environmental management test cases (e.g. Chen &

Chang 1998; Erickson *et al.* 2002; Seely *et al.* 2004; Xevi & Khan 2005; Chen *et al.* 2006; Higgs 2006; Linkov *et al.* 2006; Stirn 2006), including water management, contaminated sediments, location of waste facilities, air quality monitoring networks and forest management. Environmental management is a multi-objective problem because there are typically a number of objectives to be optimized and there are possible management actions that can be implemented; the potential effect of the management action is linked to the objectives through a model that quantifies the consequences of alternative actions. The challenge is in evaluating the performance of the action relative to the multiple objectives.

*Correspondence author. University of Washington, Quantitative Ecology and Resource Management. Box 352100 Seattle, WA 98195–2100, USA. Fax: 206 5433254. E-mail: mkenn@u.washington.edu

In many studies, the analysts solicit preferences from decision-makers that indicate the relative importance of each of the multiple objectives. Weightings for each of these are estimated from the preferences, and used to convert the objectives into a scalar-valued function that is optimized (Kim & Smith 2005). While this reduces the optimization problem and eases computation of the optimal solution, it also introduces a form of uncertainty in the optimal solution due to decisions required to quantify appropriate weights (Schoemaker & Waid 1982). The relative importance of the objectives may vary with which solicitation method is used, which decision-makers are asked, and even when preferences are solicited from a decision-maker at different times. Furthermore, the ranking of preferences can be sensitive to the weighting values (Hyde *et al.* 2005), producing variations in the preferred solution with small changes in the weights.

The idea that there is a single solution to a multi-objective optimization problem is a fallacy of the weighted sum approach. This single solution does not exist because preferences can change and are not certain themselves, and there is no single answer that minimizes all of the objectives simultaneously. A method that yields optimal solutions regardless of weights is preferable. We present a method for multi-objective optimization based on approximating the non-dominated Pareto frontier (Cohon 1978; Fig. 1; Table 1) of decision variables. The Italian economist Vilfredo Pareto, studying economic efficiency and utility, originated the concept of Pareto optimality (Cirillo 1979). It is used extensively in economics, and has been adapted for engineering and design (e.g. Statnikov & Matusov 1995). Pareto optimality has been used more recently, but sparsely, in ecological model assessment (Reynolds & Ford 1999) and in a study of optimal foraging (Rothley *et al.* 1997).

The non-dominated Pareto frontier specifies the groups of decision variable values (the optimal set) that optimize the management objectives through simultaneous optimization of a vector-valued objective function. This method provides decision-makers with a range of multiple optimal alternatives *before* the relative importance of the objectives are specified,

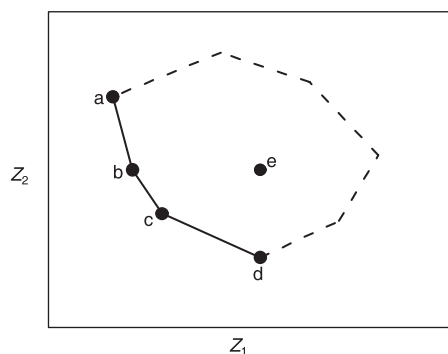


Fig. 1. Non-dominance example for the minimization of 2 objectives (Z_1 , Z_2). The dashed line encloses the feasible space. Any solutions along the solid (a,b,c,d) line belong to the non-dominated Pareto frontier, where any improvement in Z_1 occurs at a cost to Z_2 , and vice versa. The value (e) is an example of a dominated solution in the feasible space.

Table 1. 4-objective non-dominance example for FUELSOLVE case study. Each row lists objective values for unique fuels treatment distributions; the optimization problem is to minimize the objectives. In this set of 5, Fuels treatment 1–4 are not dominated by any of the other 5 results. Although in terms of dominance Fuels treatment 5 is indistinguishable from 3 and 4, it is dominated by Fuels treatment 1 and 2 (all four objectives are greater for 5 than for 1 and 2), and is not part of the non-dominated set

	LSR	Owl circle	Owl Core	Area treated
Fuels treatment 1	24.80	34.68	0.28	4518.18
Fuels treatment 2	1.08	86.56	1.76	5913.00
Fuels treatment 3	0.36	243.24	39.08	6894.72
Fuels treatment 4	53.60	22.64	0.10	4842.18
Fuels treatment 5	30.24	350.69	35.16	7864.00

thus removing the consequences of uncertainty in the weighting of objectives. It preserves the role of the decision-maker in setting preferences and re-examining both them and the proposed management model. Two major challenges can prohibit the exploration of a full multi-objective optimization problem: (a) an efficient procedure to approximate the non-dominated Pareto frontier, and (b) effective exploration of the optimization results to allow for further analysis by decision-makers. We demonstrate how an evolutionary computation algorithm can approximate the non-dominated Pareto frontier and give a four-stage process for summarizing and presenting results.

MULTI-OBJECTIVE OPTIMIZATION USING PARETO OPTIMALITY

Multi-objective optimization using the concept of non-dominance (Cohon 1978) requires approximation of the Pareto frontier, i.e. the set of all non-dominated solutions. A solution is defined to be non-dominated if there exists no other feasible solution that will give an improvement in one objective without a subsequent degradation in at least one other objective (Cohon 1978; p. 70; Fig. 1; Table 1). The optimization can be formally presented as follows:

maximize (or minimize)

$$Z(x_1, x_2, \dots, x_n) = [Z_1(x_1, x_2, \dots, x_n), Z_2(x_1, x_2, \dots, x_n), \dots, Z_p(x_1, x_2, \dots, x_n)] \quad \text{eqn 1}$$

$$X = \{(x_1, \dots, x_n \mid x_j \in \mathbb{R}); j = 1, 2, \dots, n; \quad \text{eqn 2}$$

n decision variables (X), p objectives (Cohon 1978).

In the optimization method, vector objective values are calculated for feasible combinations of values for the decision variables (x_j), and then evaluated for the relative dominance status of each vector of objectives (Z). If the entire feasible space is evaluated and objective values are deterministic then the entire non-dominated Pareto frontier is calculated (e.g. Fig. 1). With the exception of rare cases and given the size

of environmental management problems, calculation of the entire feasible space to determine the non-dominated Pareto frontier is computationally prohibitive. Search processes such as evolutionary algorithms are utilized to converge to an approximation of the non-dominated Pareto frontier; these techniques are used extensively in engineering and economics (Deb 2001).

The general optimization program can be conducted in the following steps:

1. Define and quantify the objectives.
2. Define the optimization decision variables.
3. Integrate search algorithm with the calculation of objectives.
4. Present results for post-processing.

We illustrate this method with a forest management example for the distribution of fuels reduction treatments in a watershed in the eastern Cascade Mountains, Washington State, USA, with the goal to reduce possible wildfire impact to multiple ecological values. This is a spatially complex environmental management optimization problem where management has multiple objectives.

BACKGROUND TO THE CASE STUDY

A century of fire exclusion, grazing, and selective removal of large, fire-tolerant trees in the dry forests of eastern Washington has resulted in forests of much higher density that are more prone to stand-replacing wildfires than those which occurred historically in the area (Arno & Brown 1991; Agee 1993; Wright & Agee 2004). Yet, these now dense, multi-canopied forests also serve biodiversity goals, including habitat for endangered species (Lehmkuhl *et al.* 2007). Federal legislation and regulation has defined both of these competing objectives (fuels reduction and habitat protection) without reference to the other.

The National Fire Plan, Healthy Forests Initiative & Healthy Forest Restoration Act of 2003 directs United States Federal agencies to treat these fire-prone forests by thinning and prescribed fire to reduce the risk of stand-destroying wildfire (O'Laughlin 2005). These documents all call for prioritizing hazardous fuels reduction, with an emphasis on rehabilitation and restoration of impacted forests. Yet the recovery plan for the northern spotted owl (and many other species defined as Federally threatened under the Endangered Species Act), known as the North west Forest Plan, provides protection for the owl through two mechanisms: an unmanaged buffer around each known nest site, and large protected areas known as Late Successional Reserves (LSRs) where management direction is to provide old growth, late successional conditions. This is in conflict with efforts to reduce fuels in the same landscape.

In addition, the structure of fuel treatments and their spatial distribution in a landscape are not well understood (Agee & Skinner 2005), and the effect of fuel treatments on landscape types of high ecological value (e.g. wildlife habitats and populations, old-growth forests; Huntzinger 2003; Lee & Irwin 2005) must also be considered and balanced against the risk

of losing such landscape types to catastrophic fire. Policies recognize the imperative to consider ecological consequences of fuels reduction programs and subsequent fire risk (Franklin & Agee 2003). What choice should be made about which section of a forest is to be treated by thinning and prescribed fire in order to reduce the potential spread of a wildfire? The choice requires distributing the minimum area of treatment that will minimize possible fire damage to areas of special interest, but there is no single distribution that can achieve these minima simultaneously.

This problem is similar to many environmental management problems in that it has multiple objectives that are spatially complex and potentially conflicting. Furthermore, the context of the decision is framed by multiple policy documents and the objectives described by these documents are quantified in different currencies and are not amenable to simple addition or other combinations through weightings.

FUELSOLVE CASE STUDY

The case study is conducted through the US Forest Service project called FUELSOLVE, intended to integrate ecological values into the fuels and fire management decision-making process. We chose the 23162 ha Mission Creek watershed in the Okanogan and Wenatchee National Forests of Washington State to provide a real-world landscape for analysis. The first set of objectives this project considers are the ecological values (owl habitat, late successional reserves), whereas the second set of objectives measure the cost of treatment (area assigned treatment). In the next section, the method for the four stages are described, and the results given for the post-processing example.

Methods

MULTI-OBJECTIVE OPTIMIZATION OF A RESOURCE MANAGEMENT PROBLEM

1. Define the objectives

(a) *Landscape types.* Owl activity centres have been recorded for the Mission Creek Watershed, reflecting observed activity in the years 2000–04. These centres are surrounded by two buffers of increasing size: cores (radius = 1127 m around activity centre) and circles (radius = 2931 m around activity centre); cores are given special protection in the region. LSRs are areas designated as old-growth or set aside as potential old-growth, and the study area includes both LSRs and owl activity centres (Fig. 2), with some overlap between landscape types. The fire impact objectives have effect on owl circles, owl cores and LSRs. Cost of treatment is included as an objective because an optimal solution for reducing fire spread would be to treat the entire study area, but this solution would have negative ecological and economic consequences.

(b) *Quantify objectives.* The fire impact objectives are estimated by the results of fire spread simulated on each treated landscape.

Requirements for fire spread simulations

To run the fire spread models, multiple GIS layers are required (Table 2) and the landscape data were gathered from multiple

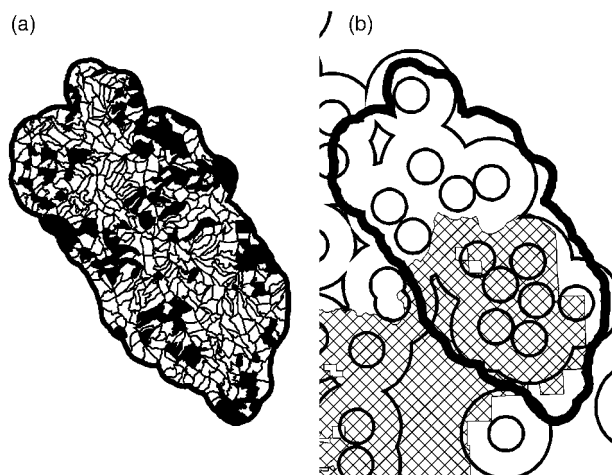


Fig. 2. (a) Study area and window for the Mission Creek Watershed. White polygons are the discrete landscape areas available for treatment. Black polygons are not available for treatment. (b) Valuable landscape types used as objectives in the optimization. The smaller circles are owl cores (1127 m buffer), the larger circles are owl circles (2931 m buffer). The shaded grid near the bottom is Late Succession Reserve (LSR).

Table 2. Landscape attributes modified for fuel treatment

Attribute	Fuels treatment value
Surface fuel model	34 (TL1)
Canopy cover	50%
Canopy height	30 m
Canopy base height	5.4 m
Canopy bulk density	0.03 kg m ⁻³

sources using ArcMap (ESRI 2004). Using the 40-model classification system (Scott & Burgan 2005), surface fuel model types were assigned to the Mission Creek watershed with shrub cover, elevation, aspect, and canopy cover. Weather files were collected from the Swauk weather station during a period from 1994 to 2003, and represented by dates ranging from 15 July to 31 August (the fire season). Only extreme weather conditions were used, under which wildfires escape initial attack and grow rapidly in size, in conjunction with a constant wind speed (10 mph) and wind direction (south-west). A study area was defined for the watershed (23162.76 ha) and a rectangular window was drawn around the study area to be included in the fire spread simulations (Fig. 2). The landscape is reduced to an ASCII grid file with 277 rows and 215 columns. Each cell in the grid represents a 90 m × 90 m pixel that contains the landscape attribute and each attribute has a separate file. The grid is orientated such that the cell (1,1) is the north-east corner of the landscape. In addition, the fire spread simulations require a particular landscape file format (*.lcp) and this was generated for the base Mission Creek Landscape (MissionCreek.lcp).

For the calculation of fire spread we utilize the command-line version of the gui program *flammap* (Finney 2002, 2003), which has two subprograms: *flammap* calculates the fire spread variables, whereas *randig* calculates realized spread based on random ignition points. The fire spread simulations generate k ignitions placed

randomly on the landscape. In the algorithm the minimum travel time of the fire to a pixel is calculated, and if that time is less than the specified fire duration the fire is determined to have spread to that pixel. Each of the k ignitions is spread independently on the unburned treated landscape, so that the fires do not interact. Given that the actual location of ignitions is difficult to predict, the goal is to characterize possible fire behaviour through multiple ignitions assigned randomly across the landscape. The more ignitions simulated, the greater the coverage of the landscape; this, however, is at great computational cost. This requires that the objective function incorporates the multiple ignitions in a manner that reduces the variability in the objective function, without undue computational burden. This is evaluated below.

To determine an appropriate fire duration for the optimization problem, preliminary fires were spread on the untreated landscape; at 4.5 days (6480 min) the mean area burned was 7500 ha. This is a reasonable size relative to other major fires in the area and all simulations were conducted with duration of 6480 min.

Fire impact

The fire spread simulation produces new landscape files whose attributes are the arrival time of the fire and the proportion (q) of the k fires that reach each pixel in the landscape (q_{rl} ; $r = 1, 2, \dots, 277$; $l = 1, 2, \dots, 215$; $q = \{0, 1/k, 2/k, \dots, 1\}$). For the fire impact objective function, the q -value is summed across all pixels that are also of a target landscape type (*fire_sum*, e.g. sum the q -values for all pixels that are in LSR). This is divided by k for a per ignition value.

$$\text{fire_sum: } Z_p(x_1, x_2, \dots, x_n) = \frac{1}{k} \sum_{r,l,p} q_{r,l,p} \quad \text{eqn 3}$$

for $p = \{1, 2, 3\}$, r_p indexes the rows relevant to objective p , and l_p indexes the columns relevant to objective p in the landscape grid file, and n is the number of units searched for treatment. The r and l values give the coordinates for each of the objective areas in the study landscape.

To determine the value for k , the variation in the per ignition *fire_sum* value for fires simulated on the untreated landscape declines steeply between 3 and 5 ignitions, then levels off at ignitions greater than 5 (Fig. 3). Given the computational cost of simulating multiple ignitions, $k = 5$ ignitions were used for all optimizations.

The cost objective (area treated) is the value of area treated under the current treatment alternative:

$$Z_p(x_1, x_2, \dots, x_n) = 0.81 \sum_{m=1}^n \sum_{r_m, l_m} I_{r_m, l_m} \quad \text{eqn 4}$$

$$I_{r_m, l_m} = \begin{cases} 1 & \text{if } x_m = 1 \text{ and } \text{pixel}_{r_m, l_m} \text{ is treated} \\ 0 & \text{otherwise} \end{cases}$$

where r_m, l_m index the row and columns for the pixels of treatable unit x_m , and 0.81 is the hectare per pixel. This can be calculated for the total area, or for the area treated of a specific landscape type. In order to demonstrate the method, we chose the simplest representation of cost of treatment as area treated. Actual cost has more complex components that depend, for example, on contiguous size, location and accessibility. When quantified these components could be included in the search.

2. Define the optimization decision variables

In the case study, there is one type of decision variable, whether or not an area of forest should be treated to reduce the potential

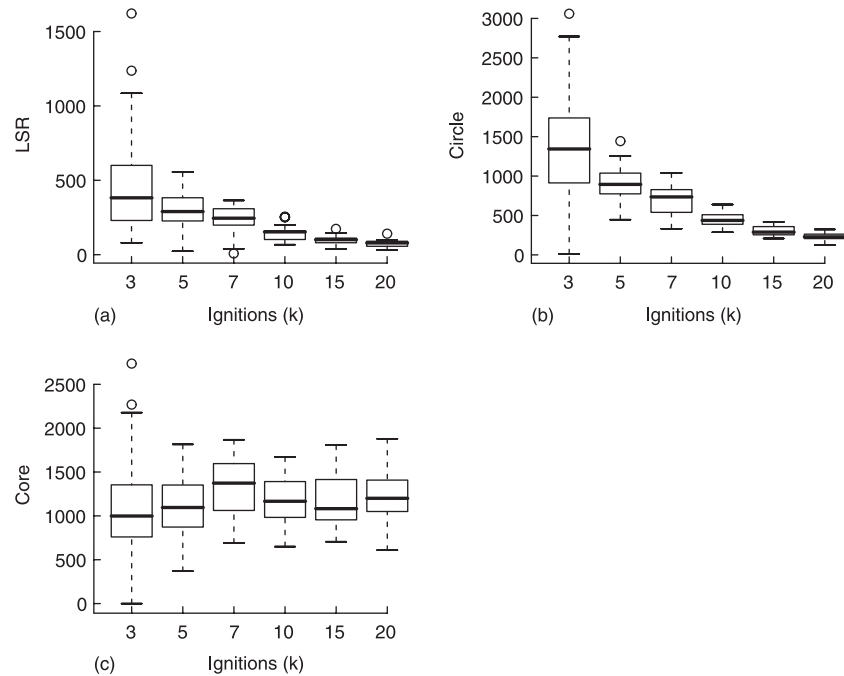


Fig. 3. Variation in the per ignition value of *fire_sum* with increasing number of ignitions for (a) LSR, (b) owl circle, and (c) owl core. Twenty simulations were conducted for each number of ignitions.

spread of fire. The study area was divided into 509 polygons of similar vegetation attributes to serve as discrete treatable areas (Fig. 2a; polygon size ranged from 19 to 95 ha). Of these, if > 30% of the pixels in a polygon contained > 40% canopy cover, that polygon was considered available to treatment. This left 413 polygons to be searched for whether treatment is assigned; each decision variable x_j ($j = 1, 2, \dots, 413$) in the optimization problem is a single potentially treated polygon.

The single fuel treatment is defined by five landscape attributes of a thin and burn fuel treatment (Agee & Skinner 2005; Table 2). In the model, if a treatable polygon is assigned fuel treatment, then all pixels in the polygon are modified to match the treatment attributes (Table 2). The possible values for the decision variables are:

$$x_j = \{0, 1\} \text{ for } j = 1, 2, \dots, 413, \text{ where } 0 = \text{no treatment,} \\ 1 = \text{assigned treatment.} \quad \text{eqn 5}$$

The optimization algorithm searches combinations of fuel treatment assignments that yield the spatial distribution of treatments that minimize the fire impact and treatment cost (area).

3. Integrate search algorithm with the calculation of objectives

The program used to solve the multi-objective problem is a multi-objective evolutionary algorithm, *Pareto_evolve* (see Appendix S1 in Supplementary material), first developed for ecological process model assessment by Reynolds & Ford (1999). The optimization algorithm initializes by randomly generating a large number of decision variable combinations (i.e. spatial allocation of treatments). There are typically 100 individual decision variable combinations, and this set is called a **population**. Once initialized, the algorithm has two major stages: evaluation and breeding.

In the evaluation stage, the effectiveness of each individual in the current population is evaluated for how well it achieves the optimization objectives. The individuals are ranked, where the non-dominated individuals of the current population are assigned Rank 1. They are assigned a fitness based on their non-dominated

ranking and a measure that reduces the fitness of individuals that are similar and increases the fitness of unique individuals, and individuals are then chosen randomly by their fitness to enter the breeding stage of the algorithm.

The individuals chosen to enter the breeding stage are called **parents** of the next population that will be evaluated. The next population is produced either through **mutation** of the parent vector (small changes in randomly chosen decision variable values, e.g. change from a zero to a 1 or vice versa), or through **cross-over** between two vectors (decision variable values are exchanged between the two parents, e.g. either the variable value doesn't change or it switches to or from 0 and 1, but it retains neighbouring values in the vector). This new population of individuals then enters the evaluation stage. Each cycle of evaluation and breeding is termed a **generation**. The algorithm is terminated either when it reaches a specified maximum generation or it converges to a unique optimum.

Pareto_evolve is a generic optimization program that the user modifies to coordinate with the user-supplied evaluation code (i.e. the fire spread simulations). It has been utilized to assess various process-based models, including stand development (Reynolds & Ford 1999), competition (Turley 2001) and shoot extension (Komuro *et al.* 2006). To configure the algorithm for a specific optimization problem the user must define the decision variable search ranges (e.g. 0, 1; step 2 above), and objective target values (e.g. 0 for minimization; step 1 above). The final output of the search is the approximated non-dominated Pareto frontier, which is presented in step 4.

4. Presentation of results for post-processing

In order to evaluate the effectiveness of the algorithm and to compare the results of different objective function combinations, three searches were conducted (Table 3). For each search the population size is 100 individuals for each generation. Given the constraints on computing power, the most generations conducted for a search is 400.

In the first search, 6 objective functions were optimized simultaneously (Search 1): fire impact on LSR, owl core and owl circle and area treated in the total study area, in LSR and in owl core. This

Table 3. Details of optimization searches and sums used to evaluate progress. The searches were conducted on a machine dedicated solely to the optimization

	Search 1	Search 2	Search 3
Population size	100	100	100
Generations	240	400	400
Ignitions per treatment (<i>k</i>)	5	5	5
Number of objectives (<i>p</i>)	6	4	2
Fire impact objectives	LSR, Owl circle, Owl core	LSR, Owl circle, Owl core	Total area
Area treated cost objectives	Total area, LSR, Owl core	Total area	Total area
Time for search	11.25 days	22.25 days	22.25 days
	Sum 1	Sum 2	
Objectives	LSR, Owl circle, Owl core, Total area treated	LSR, Owl circle, Owl core, Percentage area treated	

search represents the broadest scope of objective functions. In Search 2, all three fire impact objectives were retained, but the cost objective was restricted to total area treated (4 objectives). This focuses the optimization on the fire impact objectives, while reducing cost to a single measure. The final search is a control for the necessity of spatially explicit multiple objectives. In Search 3 there are 2 objective functions: fire impact in the total area and the total area treated. For the resulting non-dominated set, objective function values for the other fire impact objectives were calculated and compared to the first two searches.

Post-processing is conducted in three major steps that are presented in the results: (a) evaluate search progress; (b) summarize the optimal set decision space; (c) evaluate objective trade-offs in the Pareto frontier.

Search progress (a) is evaluated by calculating the minimum sum of the objectives every 20 generations. If the algorithm is making effective progress then this value should decrease over generations, and levelling off in this value indicates that the algorithm has converged. This is a coarse measurement due to the issues in weighting objectives already discussed, and particularly when considering the scale of objective values. Two sums were calculated: straight sum of four objective values (defined for Search 2; Sum 1) and the sum of the fire impact objectives with area treated expressed as a percentage (Sum 2; Table 3).

The optimal set decision space and Pareto frontier objective space are summarized for Search 2 results. The solution set is the archived set of non-dominated solutions throughout the search. The decision space (b) is summarized by the distribution of decision variable values, and by the evaluation of any dependencies among decision variables. In this example, the distribution of values is expressed as the distribution of times a polygon is assigned treatment in the optimal set. Pair-wise correlations in treatment assignment were also calculated. A strong positive correlation between polygons *a* and *b* implies that if *a* is assigned treatment, so is *b* and vice versa. A strong negative correlation between *a* and *b* implies that if *a* is assigned treatment, *b* is not.

The objective space can be summarized (c) by evaluating pair-wise relationships among objectives in the Pareto frontier, and suitable partitions of the Pareto frontier. In our example we first convert the *fire_sum* values to conditional burn probabilities (Finney 2002) by dividing by the total number of pixels for each landscape type; total area treated is presented as percentage area treated. Scatter plots of the pair-wise relationships among objective values are produced to illustrate trade-offs in the objective space. A goal of the optimization

analysis is to present informative partitions of the subspace for the purpose of identifying solutions that decision-makers will use for more detailed analysis. Percentage area treated is a convenient measure that managers use to assess the size of the fuels project; to identify solutions the objective space is partitioned into subsets of percentage area treated (area treated = 20%; 20% < area treated = 30%; 30% < area treated = 40%) and the objective trade-offs are visualized within the partitions. Individual solutions may then be chosen based on decision-maker preference.

Results

EVALUATE SEARCH PROGRESS

The number of solutions in the historical non-dominated Pareto frontier decreased as the dimensions of the objective space decreased (400-50-11 for Searches 1-3, respectively; Table 3; Fig. 4c). The value of Sum 1 tends to decrease as the search progresses (Fig. 4a), although the value for Search 1 is much higher than Searches 2 and 3. Due to computational constraints, the apparent lack of progress with respect to area treated and the large number of solutions, Search 1 was cut-off at 240 generations. The value of Sum 1 is dominated by the area treated objective, as its scale (in the thousands) is an order of magnitude larger than the scales of the fire impact objectives (in the tens and hundreds). Poor performance for this sum is explained by poor performance with respect to minimum total area treated, and this dwarfs performance with respect to fire impact. In Sum 2, area treated is calculated as percentage area treated and the scales of the objectives are more equivalent. For Sum 2, Searches 1 and 2 perform similarly as the solution evolves, and perform better than for Search 3 (Fig. 4b). These results show that Search 1 does not perform as well with respect to minimizing area treated (Fig. 4a), whereas Search 3 does not perform well with respect to the fire impact objectives (Fig. 4b). Therefore, when the spatially explicit objectives are not optimized separately (Search 3), fire impact is not reduced on the ecologically valuable landscape types.

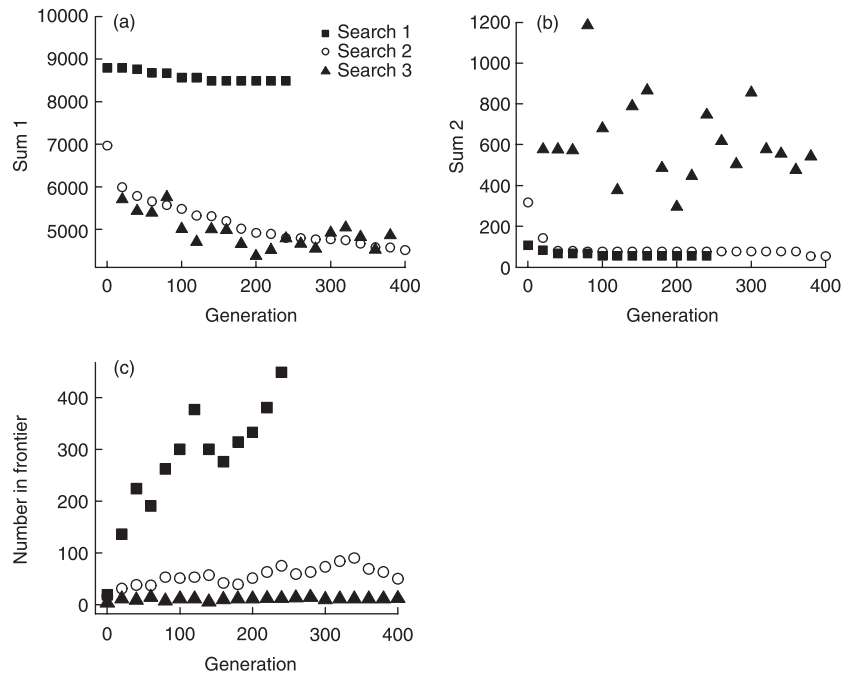


Fig. 4. Comparison of search progress for the four main objectives. (a) the minimum sum of all four objectives (Sum 1); (b) the minimum sum of the fire impact objectives with percentage area treated (Sum 2); (c) number in Pareto frontier for successive generations.

SUMMARIZE THE PARETO FRONTIER DECISION SPACE (SEARCH 2)

The distribution of the proportion of times, in the optimal set, a polygon is assigned treatment in Search 2 is bimodal with modes near 0.65 and 0.35 (Fig. 5a). For example, of the 413 polygons, 166 are assigned treatment in between 0.30 and 0.40 of the solutions; 146 are assigned treatment between 0.60 and 0.70 of the solutions. The other two searches had similar patterns in the distribution of time a polygon is assigned treatment (not pictured).

The distribution of pair-wise correlations in treatment assignment among polygons is bimodal at ± 1 (Fig. 6), a pattern that emerges early in the search evolution. This indicates strong consistent correlation among polygon treatment assignments, and these correlations must be investigated by decision-makers when considering which polygons to treat.

EVALUATE OBJECTIVE TRADE-OFFS IN THE PARETO FRONTIER

For low values of LSR (< 0.02) there are inverse relationships between burn probability of LSR (< 0.02) and burn probability of both owl core and owl circle (Search 2; Fig. 7b,d). There is a clear increasing relationship between burn probability of owl core and owl circle (Fig. 7f), and inverse relationships between percentage area treated and all three of the fire impact objectives (Fig. 7a,c,e).

Within the $\leq 20\%$ area treated partition (24 solutions), performances on the burn probabilities are highly sporadic (Fig. 8a). At the next two levels ($20\% < \text{treated area} \leq 30\%$, 11 solutions; $30\% < \text{treated area} \leq 40\%$, 14 solutions) there is more consistent performance in the burn probabilities, although there is an evident trade-off between protecting LSR and owl core (Fig. 8b,c). This trade-off is probably due

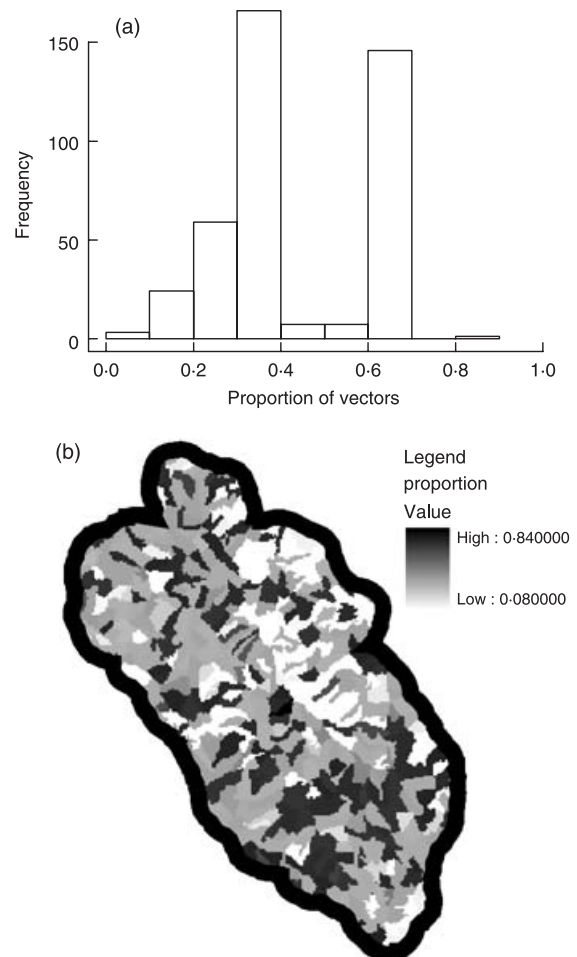


Fig. 5. (a) Distribution of the proportion of times a polygon is assigned treatment in the Search 2 non-dominated Pareto frontier. (b) Map of the spatial distribution of the proportion of times an area is assigned treatment.

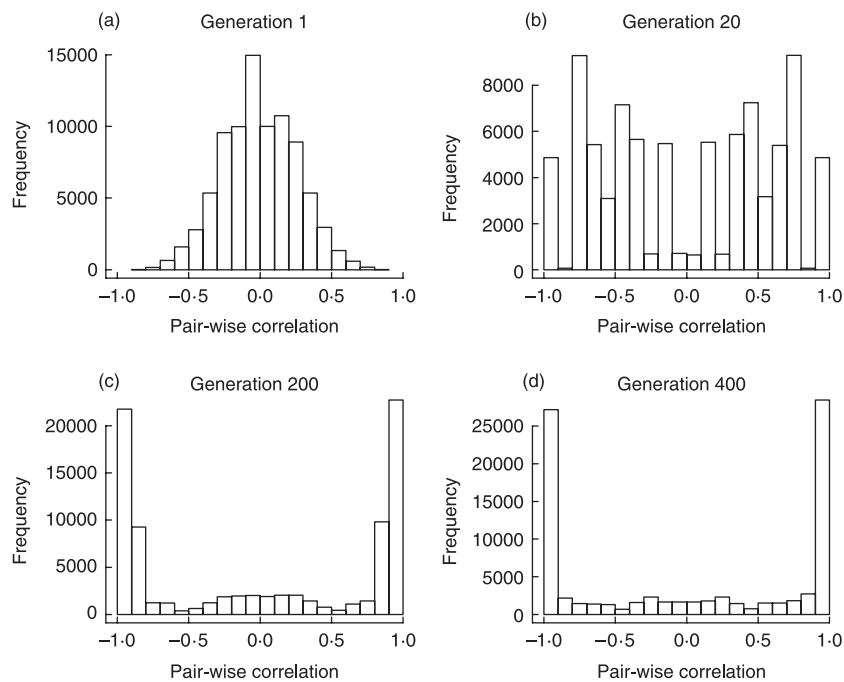


Fig. 6. Distribution of pair-wise correlations in treatment assignment among polygons in the Search 2 Generation 400 Pareto frontier.

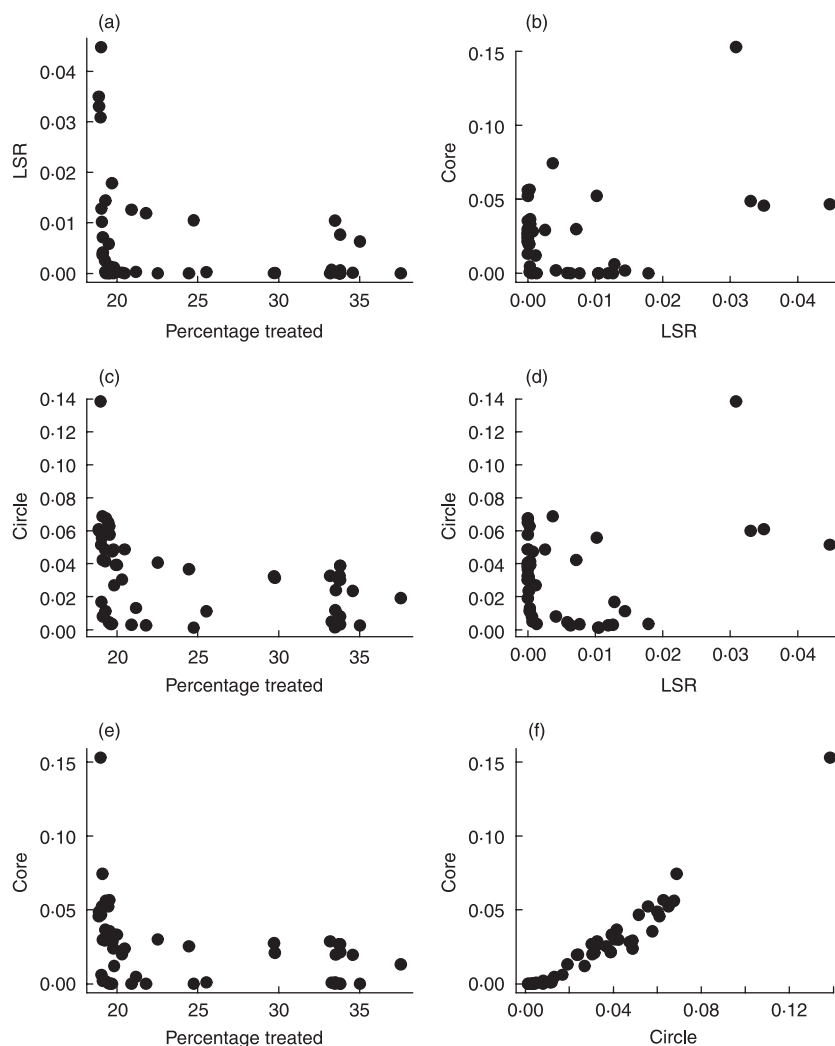


Fig. 7. Pair-wise scatter plots of objective values for the Search 2 Generation 400 Pareto frontier. All objective pairs tend to be inversely related, with the exception of a positive relationship between owl core and owl circle. There is also a potentially positive relationship between LSR and owl circle and core, respectively (but see Fig. 8).

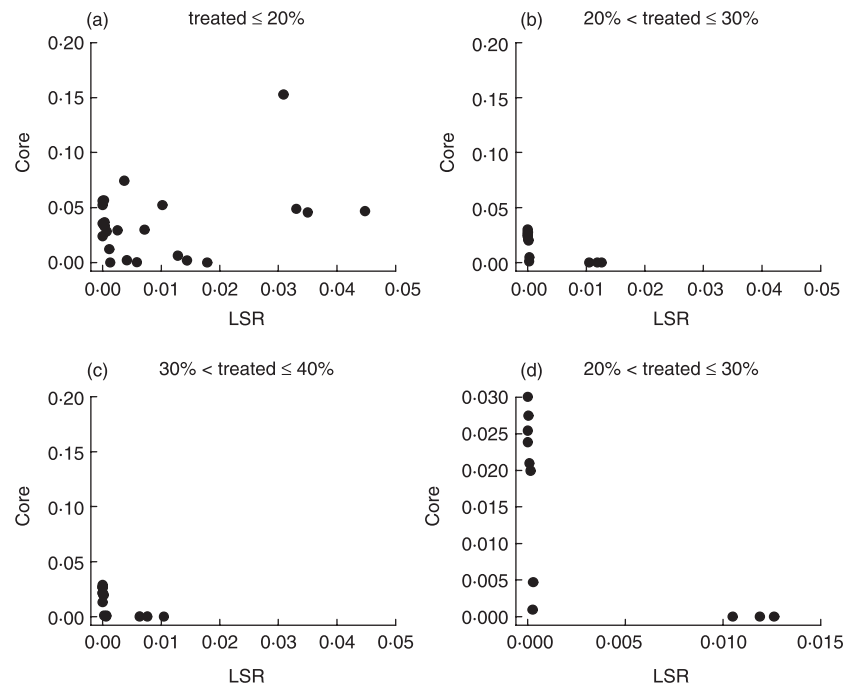


Fig. 8. Relationship between burn probabilities for LSR and Owl Core at three partitions of the percentage of area treated in the Search 2 Generation 400 Pareto frontier: (a) 20%; (b) 20% < area = 30%; (c) 30% < area = 40%. The trade-off between protecting owl core and LSR is clear from values > 20% area treated, and the effectiveness of the treatments is much more consistent when the area treated is > 20%. (d) Zoomed in plot of LSR vs. owl core for 20% < area ≤ 30% and LSR < 10, clearly showing the inverse relationship at low values of LSR.

to the way in which the treatments are distributed relative to the locations of the target areas. Since the LSR is an almost contiguous block of landscape in the bottom portion of the study area, then treatments concentrated around and within that block can protect the LSR at the expense of owl cores located above that area. Treatments distributed throughout the study area are more likely to protect more owl cores.

Discussion

MULTI-OBJECTIVE OPTIMIZATION AND ENVIRONMENTAL MANAGEMENT

A multi-objective optimization method that approximates the non-dominated Pareto frontier eliminates the need for a *priori* specification of objective weights and provides a common ground for dialogue leading to an informed compromise and decision in environmental management problems. In a weighting method there is a single solution given a set of weights, or a limited set of solutions with a sensitivity analysis on the weights. Decision-makers are not exposed to the emergent consequences of a particular set of preferences, or alternative preference sets. The method that utilizes Pareto optimality allows decision-makers to identify solutions based on the consequences of the whole set of relative preferences.

The results of the first stage of the FUELSOLVE project have resulted in a dialogue between decision-makers regarding the structure of the fuels treatment solutions and the importance of the objectives evaluated in the program. Major points of discussion have been the size of solution/optimal space and how to choose solutions within that space. The managers want to further evaluate solutions for the treatment

effect on owl habitat before the landscape is burned; such an analysis requires detailed effort that is not feasible over the entire solution set. This requires effective presentation and partitioning of the optimal space in order to choose desirable combinations. This method also gives guidelines for treatment thresholds to achieve meaningful effects for the burn objectives. The key consequence of the multi-objective optimization method that utilizes Pareto optimality is that there is not a single answer, which forces decision-makers into an exploratory investigation of the full set of solutions and to choose solutions for further analysis with the most information possible in the system.

The identification of multiple solutions allows the decision-maker to observe patterns in the distribution of decision variables, and how those might be related to objective performance. For example, the correlation of polygon treatment assignment emerged only after the searches were conducted and in the objective space there are trade-offs between the spatially explicit fire impact objectives (owl core and LSR).

FUELSOLVE: decision space

In the Search 2 historical Pareto frontier some polygons are assigned treatment more consistently than others (Fig. 5a), and strong positive and negative correlations emerge in treatment assignment among polygons (Fig. 6d). The post-processing analysis shows where highly correlated polygons could be combined in order to reduce the complexity of the decision search space. This would require additional criteria, such as vegetation-type and polygon contiguity, in order to choose which polygons are combined. These polygons can be evaluated for which landscape characteristics are associated

with fuel treatment assignment, and in this way Pareto frontier calculation aids a progressive analysis of a management optimization problem.

A reduced decision space with larger individual polygons may result in a more efficient optimization search given fewer possible search alternatives. This would also align the individual areas treated more closely with those that might be chosen based on manager experience without an optimization analysis (larger individual sections of landscape treated). These possibilities will be explored in more detail in the next stage of the FUELSOLVE project.

FUELSOLVE: objective space

Comparison of the three searches demonstrates that effective optimization of spatially complex objectives requires simultaneously minimizing the objectives with a spatial component (Fig. 4a,b). Inclusion of particular landscape types for area treated only increases the size of the non-dominated Pareto frontier, at the detriment to total area treated (Fig. 4c). This may be a result of the higher dimensions of that objective space (6 rather than 4 objectives). The method of multi-objective optimization with Pareto optimality demonstrates that the relative importance of the objectives, and which objectives are calculated, both have apparent consequences for the kind of solutions that are found through optimization. Furthermore, it is also clear that if only one of the scalar-valued functions represented by Sum 1 and Sum 2 were optimized, a different optimal solution would be identified.

We chose relatively simple, easily quantifiable objectives (eqns 3 and 4), but an advantage of this multi-objective method is that it can incorporate additional objectives, or eliminate those concluded to be less informative than originally believed (reduction from 6 to 4 objectives between Search 1 and Search 2). This is accomplished without the need to find a common currency for the objectives, and without a complicated procedure to reassign the relative weighting of each objective. For the optimization of fuels treatment distribution it may become a priority to quantify treatment and fire effects on areas that have more direct human impact such as the wildland urban interface (wui; Berry & Hessel 2004). The cost of treatment may require a more sophisticated calculation that incorporates issues like the size of the treated area, topography (slope and elevation), timing and fuels type (Rideout & Omi 1995; Berry & Hessel 2004).

With respect to computation effort, it is possible to get an approximation of the archived non-dominated frontier when a new objective is added by taking the current set, calculating the new objective and re-evaluating dominance. This method, however, is unlikely to discover the lowest values of the new objective. In an opposite scenario, one can take the results of the 6-objective Search 1 and evaluate the non-dominated frontier with respect to the four objectives of Search 2. In this case, results similar to the full optimization Search 2 are yielded; this greatly reduces the computational burden with respect to conducting the full analysis, and is a feasible alternative to conducting the search when the objective space is reduced.

For the current objective space, area treated is a convenient benchmark to produce smaller partitions of the Pareto frontier for further postprocessing. Identification of relevant partitions of the objective space provides decision-makers with the consequences of choosing different sets of relative importance for the objectives given the emergent objective trade-offs. Overall, there is an inverse relationship between the percentage area treated and burn probability in all three of the landscape types, also observed by Finney (2006) and Finney *et al.* (2006). In addition, trade-offs emerge in protecting areas with distinct spatial distributions (owl cores and LSRs; Fig. 8d). Finney *et al.* (2006) found through simulation that if more than 40% of the landscape is treated, random distributions of treatments perform as well as spatially optimized distribution. This conclusion seems supported by our results in which no non-dominated solution is greater than 40% area treated.

EMERGENT CHALLENGES IN THE OPTIMIZATION PROGRAM

Computation time

It is possible, given the large decision space in this problem, that the generations conducted for the searches is insufficient for the algorithm to converge to the 'true' optimal solution. However, the progress with respect to the minimum sum of the objectives is marginal beyond the early generations, and such incremental changes may not be important from the management perspective. Furthermore, for the purpose of planning fuel treatment distributions, the computation time for the optimization searches given the current number of generations is impractical. This is the first stage in a progressive analysis, and the goal of the FUELSOLVE project is to also evaluate the impact of the treatment distribution on the ecological values of owl habitat and LSRs. Such an analysis requires both a small set of solutions, and extensive effort in evaluating landscape attributes that result from the treatment alternatives.

A first issue is the number of generations and the population size necessary to achieve an acceptable convergence in the non-dominated Pareto frontier. The minimum sum of objectives in the Pareto frontier levels off relatively early in the search (Fig. 4a,b), and the pattern of pair-wise correlations also evolves quickly (Fig. 6). It may be that the main general patterns in the treatment allocations in the landscape evolve early in the search. Furthermore, in a similar problem of utilizing an evolutionary algorithm for forest harvest scheduling, Falcão & Borges (2001) found a population size of 30 was sufficient for an efficient convergence, although they conducted their single-objective optimization for over 400 000 iterations. Their problem had more than 1000 constraints, and their decision space encompassed 696 management units, with up to 400 management alternatives. It may be that convergence is possible in the fuels treatment distribution problem with a smaller population size (currently 100), significantly reducing the total computation time and possibly allowing for more generations in the search.

The overwhelming computer effort in the optimization program is for the programs utilized to calculate fire spread (*randig* and *flammap*). These programs are multi-threaded for parallel computing, and current simulations were conducted on a dual processor Dell server that was dedicated solely to the optimization. The computation time can be further reduced by increasing the number of processors, and the evolutionary algorithm itself is an ideal candidate for parallel computing since the individuals in the population for each generation of the evolutionary algorithm can be run independently.

Selection of solutions from the non-dominated Pareto frontier

The total number of solutions in the non-dominated Pareto frontier for Search 2 (50 total unique treatment distributions) is unwieldy for detailed analysis by decision-makers. An important step in the dialogue with decision-makers is the informative reduction of the solution space to a few alternatives that can be investigated individually in detail. We have already shown that post-processing of partitions of the Pareto frontier, determined by ranges of area treated, can reduce the number of alternatives and choices that can be made among alternatives in the resulting partitions (Fig. 8). In our method, this essential step is conducted after the full range of trade-offs is visualized and allows the most informed decision of relative objective preference for the forest managers.

Reducing the dimensions of the objective space also effectively reduces the total number of alternatives in the final Pareto frontier (Fig. 4c), but the fire impact results of Search 3 show that reduction of the number of solutions is not a worthy goal in itself. The objective space should not be reduced without consideration of the consequences for other objectives.

Conclusions

Pareto optimality provides a methodology that incorporates competing objectives and provides decision-makers with multiple alternatives, each optimal for different weightings of objectives, which is important for environmental management (Kim & Smith 2005). It can also show which combinations of objectives can be optimized simultaneously, and which cannot, which can lead to changes in the management model. A full presentation of the realized trade-offs can be made to the public and utilized to garner trust in the competence of the decision-making process, a key to public acceptance (Winter *et al.* 2004).

Acknowledgements

The FUELSOLVE project is an ongoing collaboration between the University of Washington, College of Forest Resources, the Okanogan & Wenatchee National Forests, and the US Forest Service, Pacific North-west Research Station under Joint Venture Agreement PNW 03-JVA-11261927. Nick Brazee provided the initial GIS analysis. Katie Grieve, Derek McClure and Bert Loosmore and two reviewers gave helpful comments to improve the manuscript.

References

- Agee, J.K. (1993) *Fire Ecology of Pacific Northwest Forests*. Island Press, Washington DC.
- Agee, J.K. & Skinner, C.N. (2005) Basic principles of forest fuel reduction treatments. *Forest Ecology and Management*, **211**, 83–96.
- Arno, S.F. & Brown, J.K. (1991) Overcoming the paradox in managing wildland fire. *Western Wildlands*, **Spring**, 40–46.
- Berry, A.H. & Hesseln, H. (2004) The effect of the wildland–urban interface on prescribed burning costs in the Pacific Northwestern United States. *Journal of Forestry*, **102**, 33–37.
- Chen, H.W. & Chang, N.-B. (1998) Water pollution control in the river basin by fuzzy genetic algorithm-based multi-objective programming modeling. *Water Science and Technology*, **37**, 55–63.
- Chen, C.-H., Liu, W.-L. & Chen, C.-H. (2006) Development of a multi-objective planning theory and system for sustainable air quality monitoring networks. *Science of the Total Environment*, **354**, 1–19.
- Cirillo, R. (1979) *The Economics of Vilfredo Pareto*. Frank Cass Ltd, Totowa, NJ.
- Cohon, J.L. (1978) Multi-objective programming and planning. *Mathematics in Science and Engineering*, Vol. 140. Academic Press Inc., London.
- Deb, K. (2001) *Multi-Objective Optimization Using Evolutionary Algorithms*. John Wiley & Sons, Chichester.
- Erickson, M., Mayer, A. & Horn, J. (2002) Multi-objective optimal design of groundwater remediation systems: application of the niched Pareto genetic algorithm (NPGA). *Advances in Water Resources*, **25**, 51–65.
- ESRI (2004) *ArcMap 9.0*. Redlands, CA.
- Falcão, A.O. & Borges, J.G. (2001) Designing an evolution program for solving integer forest management scheduling models: an application in Portugal. *Forest Science*, **47**, 158–168.
- Finney, M.A. (2002) Fire growth using minimum travel time methods. *Canadian Journal of Forest Research*, **32**, 1420–1424.
- Finney, M.A. (2003) Calculation of fire spread rates across random landscapes. *International Journal of Wildland Fire*, **12**, 167–174.
- Finney, M.A. (2006) A computational method for optimizing fuel treatment locations. *Fuels Management: How to Measure Success. Conference Proceedings 2006, 28–30 March, Portland, OR* (eds P.L. Andrews & B.W. Butler), pp. 107–122. Proceedings RMRS-P-41. US Department of Agriculture, Forest Service, Rocky Mountain Research Station, Fort Collins, CO.
- Finney, M.A., Seli, R.C., McHugh, C.W., Ager, A.A., Bahro, B. & Agee, J.K. (2006) Simulation of long-term landscape-level fuel treatment effects on large wildfires. *Fuels Management: How to Measure Success. Conference Proceedings 2006, 28–30 March, Portland, OR* (eds P.L. Andrews & B.W. Butler), pp. 125–147. Proceedings RMRS-P-41. US Department of Agriculture, Forest Service, Rocky Mountain Research Station, Fort Collins, CO.
- Franklin, J.F. & Agee, J.K. (2003) Forging a science-based national forest fire policy. *Issues in Science and Technology*, **20**, 59–66.
- Higgs, G. (2006) Integrating multi-criteria techniques with geographical information systems in waste facility location to enhance public participation. *Waste Management and Research*, **24**, 105–117.
- Huntzinger, M. (2003) Effects of fire management practices on butterfly diversity in the forested western United States. *Biological Conservation*, **113**, 1–12.
- Hyde, K.M., Maier, H.R. & Colby, C.B. (2005) A distance-based uncertainty analysis approach to multi-criteria decision analysis for water resources decision making. *Journal of Environmental Management*, **77**, 278–290.
- Kim, K.-J. & Smith, R.L. (2005) Systematic procedure for designing processes with multiple environmental objectives. *Environmental Science and Technology*, **39**, 2394–2405.
- Komuro, R., Ford, E.D. & Reynolds, J.H. (2006) The use of multi-criteria assessment in developing a process model. *Ecological Modelling*, **197**, 320–330.
- Lee, D.C. & Irwin, L.L. (2005) Assessing risks to spotted owls for forest thinning in fire-adapted forests of the western United States. *Forest Ecology and Management*, **211**, 191–209.
- Lehmkuhl, J., Kennedy, M., Ford, E.D., Singleton, P.H., Gaines, W.L. & Lind, R.L. (2007) Seeing the forest for the fuel: Integrating ecological values and fuels management. *Forest Ecology and Management*, **246**, 73–80.
- Linkov, I., Satterstrom, K., Kiker, G., Seager, T.P., Bridges, T., Gardner, K.H., Rogers, S.H., Belluck, D.A. & Meyer, A. (2006) Multicriteria decision analysis: a comprehensive decision approach for management of contaminated sediments. *Risk Analysis*, **26**, 61–78.
- O’Laughlin, J. (2005) Policies for risk assessment in federal land and resource management decisions. *Forest Ecology and Management*, **211**, 15–27.
- Reynolds, J. & Ford, E.D. (1999) Multi-criteria assessment of ecological process models. *Ecology*, **80**, 538–553.

- Rideout, D.B. & Omi, P.N. (1995) Estimating the cost of fuels treatment. *Forest Science*, **41**, 664–674.
- Rothley, K.D., Schmitz, O.J. & Cohon, J.L. (1997) Foraging to balance conflicting demands: novel insights from grasshoppers under predation risk. *Behavioral Ecology*, **8**, 551–559.
- Schoemaker, P.J. & Waid, C.C. (1982) An experimental comparison of different approaches to determining weights in additive utility models. *Management Science*, **28**, 182–196.
- Scott, J.H. & Burgan, R.E. (2005) *Standard fire behavior fuel models: a comprehensive set for use with Rothermel's surface fire spread model*. General Technical Report RMRS-GTR-153. USDA Forest Service, Rocky Mountain Research Station, Missoula, MT.
- Seely, B., Nelson, J., Wells, R., Peter, B., Meitner, M., Anderson, A., Harshaw, H., Sheppard, S., Bunnell, F.L., Kimmins, H. & Harrison, D. (2004) The application of a hierarchical, decision-support system to evaluate multi-objective forest management strategies: a case study in northeastern British Columbia, Canada. *Forest Ecology and Management*, **199**, 283–305.
- Statnikov, R.B. & Matusov, J.B. (1995) *Multicriteria Optimizatoin and Engineering*. Chapman & Hall, New York.
- Stirn, L.Z. (2006) Integrating fuzzy analytic hierarchy process with dynamic programming for determining optimal forest management decisions. *Ecological Modelling*, **194**, 296–305.
- Turley, M.C. (2001) *Investigating alternative ecological theories using multiple criteria assessment with evolutionary computation*. PhD Dissertation, University of Washington, Seattle, WA.
- Winter, G., Vogt, C.A. & McCaffrey, S. (2004) Examining social trust in fuels management strategies. *Journal of Forestry*, **102**, 8–15.
- Wright, C.S. & Agee, J.K. (2004) Fire and vegetation history in the eastern Cascade Mountains, Washington. *Ecological Applications*, **14**, 443–459.
- Xevi, E. & Khan, S. (2005) A multi-objective optimisation approach to water management. *Journal of Environmental Management*, **77**, 269–277.

Received 6 March 2007; accepted 11 June 2007

Handling Editor: Rob Freckleton

Supplementary material

The following supplementary material is available for this article.

Appendix S1. *Pareto_evolve* source files and documentation.

This material is available as part of the online article from:

[http://www.blackwell-synergy.com/doi/full/10.1111/](http://www.blackwell-synergy.com/doi/full/10.1111/j.1365-2664.2007.01367.x)

[j.1365-2664.2007.01367.x](http://www.blackwell-synergy.com/doi/full/10.1111/j.1365-2664.2007.01367.x)

(This link will take you to the article abstract.)

Please note: Blackwell Publishing is not responsible for the content or functionality of any supplementary materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.